arXiv:1703.03769v1 [math.OC] 10 Mar 2017

A Novel Convex Relaxation for Non-Binary Discrete Tomography

Jan Kuske[†], Paul Swoboda[‡] and Stefania Petra[†]

[†]MIG, Inst. Appl. Mathematics, Heidelberg University
[‡] Institute of Science and Technology (IST), Austria

Abstract. We present a novel convex relaxation and a corresponding inference algorithm for the non-binary discrete tomography problem, that is, reconstructing discrete-valued images from few linear measurements. In contrast to state of the art approaches that split the problem into a continuous reconstruction problem for the linear measurement constraints and a discrete labeling problem to enforce discrete-valued reconstructions, we propose a joint formulation that addresses both problems simultaneously, resulting in a tighter convex relaxation. For this purpose a constrained graphical model is set up and evaluated using a novel relaxation optimized by dual decomposition. We evaluate our approach experimentally and show superior solutions both mathematically (tighter relaxation) and experimentally in comparison to previously proposed relaxations.

1 Introduction

We study the discrete tomography problem, that is reconstructing a discrete-valued image from a small number of linear measurements (tomographic projections), see Figure 1 for an illustration. The main difficulty in reconstructing the original image is that there are usually far too few measurements, making the problem ill-posed. Hence, it is common to search for a discrete-valued image that (i) satisfies the measurements and (ii) minimizes an appropriate energy function.

More generally, the discrete tomography problem can be regarded as reconstructing a *discrete-valued synthesis/analysis-sparse* signal from few measurements which is observed by *deterministic* sensors A. This is, in turn, a special instance of the compressed sensing problem [7], for which it has been shown that discreteness constraints on the possible values of the reconstructed function can significantly reduce the number of required measurements [10].

Acknowledgments: We gratefully acknowledge support by the DFG (German Science Foundation), Grant GRK 1653. This work is partially funded by the European Research Council under the European Unions 7th Framework Programme (FP7/2007-2013)/ERC grant agreement no 616160. The authors would like to thank Vladimir Kolmogorov for helpful discussions.

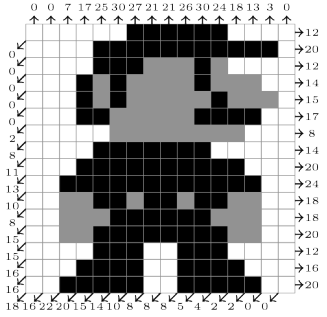

Fig. 1. Examplary discrete tomography example on three-valued image (white=0, gray=1, black=2) with three projection directions: horizontal, vertical and diagonal (right upper to lower left). Horizontal sums are written on the right, vertical sums on the top and diagonal sums on the left and bottom.

However, the discreteness constraint leads to great computational challenges. Simple outer convex relaxations coming from continuous scenarios are doomed to fail, as they will not output discrete solutions, unless the signal sparsity satisfies favourable relations and A is well-conditioned on the class of sparse signals, e.g. a random matrix. In fact, in most practical scenarios the projection matrices fall short of assumptions that underlie rigorous compressed sensing theory (like e.g. the restricted isometry property [7]), and standard algorithms from the continuous ℓ_1-setting cannot be applied any more. Rounding continuous solutions will on the other hand render the solutions infeasible for the measurement constraints.

Therefore, algorithms exploiting the *combinatorial structure* of the discrete tomography problem are necessary to successfully exploit discreteness as prior knowledge and to reduce the number of required measurements.

Related work. Several algorithms have been proposed to solve the discrete tomography problem. Among them are (i) linear programming-based algorithms [9,17], (ii) belief propagation [8], (iii) network flow techniques [3], (iv) convex-convave programming [19,14] (v) evolutionary algorithms [2] and other heuristic algorithms [4,6,11,12]. Not all approaches are applicable to the general discrete tomography problem we treat here: algorithms [9,17,8,6,11,12] only support binary labels, while [4,3] solves only the feasiblity problem and does not permit any energy. Algorithms [19,14] are applicable to the setting we propose but are purely primal algorithms and do not output dual lower bounds (which we do). Hence, one cannot judge proximity of solutions computed by [19,14] to global optimal ones, prohibiting its use in branch and bound. In case convex relaxations were considered [9,19,17,3], they were less tight than the one we propose, leading to inferior dual bounds.

Contribution. We propose the (to our knowledge) first LP-based algorithm for the non-binary discrete tomography problem. In particular, we

– recast the discrete tomography problem as a Maximum-A-Posteriori inference problem in a graphical model with additional linear constraints coming from the tomographic projections in Section 2,

- construct higher order factors in the graphical model such that a feasible solution to the higher order factors coincides with solutions feasible for the tomographic projections and
- present an efficient exact algorithm for solving the special case when exactly one ray constraint is present and the energy factorizes as a chain (such a problem will be called a one-dimensional discrete tomography problem) in Section 3,
- decompose the whole problem into such subproblems and solve this decomposition with bundle methods in Section 4.

Our approach leads to significantly tighter bounds as compared to generalising previously proposed relaxations to the non-binary case, see Proposition 1 in Section 4 and experiments in Section 5. Code and datasets are available on GitHub[1].

Notation. Let $[a, b] = \{a, \dots, b\}$ be the set of natural numbers between a and b. For $x \in \mathbb{R}$ we denote by $\lfloor x \rfloor$ and $\lceil x \rceil$ the floor and ceiling function.

2 Problem Statement

The discrete tomography problem we study consists in finding a discrete labeling $x \in \{0, 1, \dots, k-1\}^n$ such that (i) tomographic projection constraints given by $Ax = b$ with $A \in \{0,1\}^{m \times n}, b \in \mathbb{N}^m$ are fulfilled and (ii) x minimizes some *energy* $E : \{0, 1, \dots, k-1\}^n \to \mathbb{R}$. We assume that E factorizes according to a pairwise graphical model: given a graph $\mathsf{G} = (\mathsf{V}, \mathsf{E})$, together with a *label space* $\mathcal{X}_\mathsf{V} := \prod_{u \in \mathsf{V}} \mathcal{X}_u$, $\mathcal{X}_u := \{0, 1 \dots, k-1\}$ $\forall u \in \mathsf{V}$, the energy is a sum of *unary potentials* $\theta_u : \mathcal{X}_u \to \mathbb{R}$ $\forall u \in \mathsf{V}$ and *pairwise* ones $\theta_{uv} : \mathcal{X}_u \times \mathcal{X}_v \to \mathbb{R}$ $\forall uv \in \mathsf{E}$. The full problem hence reads

$$\min_{x \in \mathcal{X}_\mathsf{V}} E(x) := \sum_{u \in \mathsf{V}} \theta_u(x_u) + \sum_{uv \in \mathsf{E}} \theta_{uv}(x_u, x_v) \quad \text{s.t.} \quad Ax = b\,. \tag{1}$$

For the discrete tomography problem we usually choose G to be a grid graph corresponding to the pixels of the image to be reconstructed, zero unary potentials $\theta_u \equiv 0$ $\forall u \in \mathsf{V}$, **as no local information about the image values is known**, and pairwise potentials $\theta_{uv} = g(x_u - x_v)$ penalize intensity transitions, e.g. $g(\cdot) = |\cdot|$ (TV) or $g(\cdot) = \min(1, |\cdot|)$ (Potts). Such choice of pairwise potentials assigns small energy to labelings x with a regular spatial structure.

3 One-Dimensional Non-Binary Discrete Tomography

A natural decomposition of the discrete tomography problem (1) consists of (i) considering a subproblem for each ray constraint separately and (ii) joining them together via Lagrangian variables. We will study the first aspect below and

[1] `https://github.com/pawelswoboda/LP_MP`

the second one in the next section. In particular, let $U = \{u_1, \dots, u_n\} \subseteq \mathsf{V}$ be the variables from a single ray constraint $x_{u_1} + \ldots + x_{u_n} = b$ corresponding to a row of the projection matrix A in (1). Assume that pairwise potentials form a chain, i.e. they are $\theta_{u_i u_{i+1}}$, $i = 1, \dots, n-1$. The *one-dimensional discrete tomography problem* is

$$\begin{aligned} \min_{(x_1,\dots,x_n)\in\mathcal{X}_U} & \sum_{u\in U} \theta_u(x_u) + \sum_{i\in[n-1]} \theta_{u_i,u_{i+1}}(x_{u_i}, x_{u_{i+1}}) \\ \text{s.t.} & \sum_{u\in U} x_u = b\,. \end{aligned} \tag{2}$$

We present an exact linear programming relaxation and an efficient message-passing routine to solve (2) below.

3.1 Linear Programming Model

The one-dimensional discrete tomography subproblem (2) could naively be solved by dynamic programming by going over all variables $u_1, \dots, u_n$ sequentially. This however would entail quadratic space complexity in the number of nodes in U, as the state space for variable u_i would need to include costs for all possible labels x_{u_i} and all values of the intermediate sum $\sum_{j=1}^{i-1} x_{u_j}$. The latter sum can have $1 + \sum_{j=1}^{i-1}(|\mathcal{X}| - 1)$ possible values. To achieve a better space complexity, we will recursively (i) equipartition variables $u_1, \dots, u_n$, (ii) define LP-subproblems in terms of so-called *counting factors* which are exact on each subpartition and (iii) join them together to eventually obtain an exact LP-relaxation for (2). Our approach is inspired by [16].

Partition of variables. Given the nodes $u_1, \dots, u_n$, we choose an equipartition $\Pi_1 = \{u_1, \dots, u_{\lfloor n/2 \rfloor}\}$ and $\Pi_2 = \{u_{\lfloor n/2 \rfloor + 1}, \dots, u_n\}$. We recursively equipartition Π_1 into $\Pi_{1,1}$ and $\Pi_{1,2}$ and do likewise for Π_2. For $u_1, \dots, u_8$ we obtain a recursive partitioning as in Figure 2.

Counting factors. Given an interval, a *counting factor* holds the states of its left and right end and the value of the intermediate sum.

Definition 1 (Counting label space). *The counting label space for interval* $[i, j]$ *is* $\mathcal{X}_{i:j} := \mathcal{X}_{u_i} \times \mathcal{S}_{i:j} \times \mathcal{X}_{u_j}$ *with* $\mathcal{S}_{i:j} = \{0, 1, \dots, 1 + \sum_{l=i+1}^{j-1}(|\mathcal{X}_{u_l}| - 1)\}$ *holding all possible intermediate sums. A* counting label $x_{i:j}$ *consists of the three components* $(x_{u_i}, s_{i:j}, x_{u_j})$*: its left endpoint label* x_{u_i}*, intermediate sum* $s_{i:j} := x_{u_{i+1}} + \ldots + x_{u_{j-1}}$ *and right endpoint label* x_{u_j}*.*

See again Figure 2 for the exemplary case $U = \{u_1, \dots, u_8\}$.

For interval $[i, j]$ there are $|\mathcal{X}_{i:j}| = |\mathcal{X}_{u_i}| \cdot |\mathcal{X}_{u_j}| \cdot |\mathcal{S}_{i:j}|$ distinct counting labels. We associate to each counting factor *counting marginals* $\mu_{i:j}$ satisfying $\{\mu_{i:j} \in \mathbb{R}_+^{|\mathcal{X}_{i:j}|} : \sum_{x_{i:j}\in\mathcal{X}_{i:j}} \mu_{i:j}(x_{i:j}) = 1\}$.

Assuming an uniform label space $|\mathcal{X}_u| = k\ \forall u \in \mathsf{V}$, the total space complexity of all counting factors is $O(k^3 \cdot n \cdot \log(n))$, hence subquadratic in the number of nodes in U.

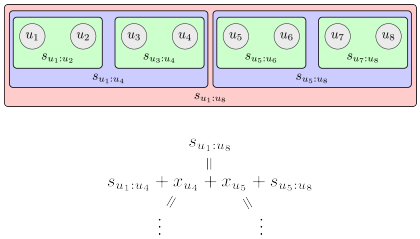

Fig. 2. *Partition of variables*: The chain of variables $U = \{u_1, \dots, u_8\}$ (red boxes) is first partitioned into $\Pi_1 = \{u_1, u_2, u_3, u_4\}, \Pi_2 = \{u_5, u_6, u_7, u_8\}$ (blue boxes), and then into $\Pi_{1,1} = \{u_1, u_2\}, \Pi_{1,2} = \{u_3, u_4\}, \Pi_{2,1} = \{u_5, u_6\}, \Pi_{2,2} = \{u_7, u_8\}$ (green boxes). *Counting sums:* The topmost sum $s_{u_1:u_8} = x_{u_2} + \ldots + x_{u_7}$ is composed of sums $s_{1:4}$ for partition Π_1 and $s_{5:8}$ for Π_2. Subsums are generated recursively for $s_{1:4}$ and $_{5:8}$ again.

Joining counting factors. Assume the partitioning of variables has produced two adjacent subsets $\Pi = \{u_i, \dots, u_j\}$ and $\Pi' = \{u_{j+1}, \dots, u_l\}$, which were constructed from their common subset $\Pi \cup \Pi' \subseteq U$. The associated three counting factors with marginals $\mu_{i:j}, \mu_{j+1:l}$ and $\mu_{i:l}$ introduced above shall be consistent with respect to each other.

Definition 2 (Label consistency). *Label $x_{i:l} \in \mathcal{X}_{i:l}$, $x_{i:j} \in \mathcal{X}_{i:j}$ and $x_{j+1:l}$ are* consistent *with each other, denoted by $x_{i:j}, x_{j+1:l} \sim x_{i:l}$ iff (i) left endpoint labels of $x_{i:j}$ and $x_{i:l}$ match, (ii) right endpoint labels of $x_{j+1:l}$ and $x_{i:l}$ match and (ii) intermediate sums match $s_{i:l} = s_{i:j} + x_{u_j} + x_{u_{j+1}} + s_{j+1:l}$.*

We enforce this by introducing a *higher order marginal* $\mu_{i:j:l} \in \mathbb{R}_+^{\mathcal{X}_{i:j} \times \mathcal{X}_{j+1:l}}$ to bind together $\mu_{i:j}, \mu_{j+1:l}$ and $\mu_{i:l}$.

$$\sum_{x_{j+1:l}} \mu_{i:j:l}(x_{i:j}, x_{j+1:l}) = \mu_{i:j}(x_{i:j}) \qquad \forall x_{i:j} \in \mathcal{X}_{i:j} \tag{3}$$

$$\sum_{x_{i:j}} \mu_{i:j:l}(x_{i:j}, x_{j+1:l}) = \mu_{j+1:l}(x_{j+1:l}) \qquad \forall x_{j+1:l} \in \mathcal{X}_{j+1:l} \tag{4}$$

$$\sum_{x_{i:j}, x_{j+1:l} \sim x_{i:l}} \mu_{i:j:l}(x_{i:j}, x_{j+1:l}) = \mu_{i:l}(x_{i:l}) \qquad \forall x_{i:l} \in \mathcal{X}_{i:l} \tag{5}$$

The recursive arrangement of counting factors is illustrated in Figure 3.

Remark 1. The constraints between $\mu_{i:j:l}$ and $\mu_{i:j}$ and $\mu_{j+1:l}$ are analoguous to the marginalization constraints between pairwise and unary marginals in the local polytope relaxation for pairwise graphical models [18]. The constraints between $\mu_{i:j:l}$ and $\mu_{i:l}$ however are different. Hence, specialized efficient solvers for inference in graphical models cannot be applied.

Costs. Above we have described the polytope for the one-dimensional discrete tomography problem (2). The LP-objective consists of vectors $\theta_{i:j}$ for each

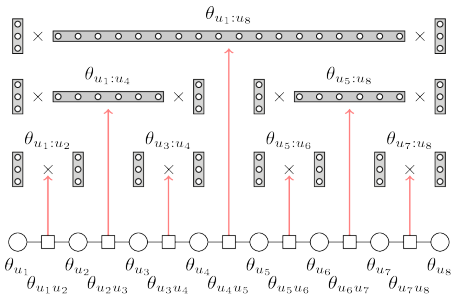

Fig. 3. Illustration of the tree construction used for solving the one-dimensional tomography problem (2) with $\ell = 3$, i.e. three labels. On the bottom is the original chain, above it are three layers of counting factors. Each counting factor keeps track of the label x_u (left vertical bar), x_v (right vertical bar) and the intermediate sum $\sum_{w \in U: u<w<v} x_w$ (middle horizontal bar).

counting marginal and $\theta_{i:j:l}$ for each higher order marginal. Accounting for the pairwise costs in (2) we set $\theta_{i:j}(x_{i:j}) := \begin{cases} \theta_{u_i u_j}(x_{u_i}, x_{u_j}), & i+1=j \\ 0, & \text{otherwise} \end{cases}$ for the counting factors and for the higher order factors we set $\theta_{i:j:l}(x_{i:j}, x_{j+1:l}) := \theta_{u_j u_{j+1}}(x_{u_j}, x_{u_{j+1}})$. For the projection constraint in (2) we set costs of the top counting marginal as $\theta_{1:n}(x_{1:n}) := \begin{cases} 0, & x_{u_1} + s_{1:n} + x_{u_n} = b \\ \infty, & \text{otherwise} \end{cases}$.

3.2 Message Passing Algorithm

Above we have introduced a linear program formulation for the one-dimensional discrete tomography problem (2). While it is possible to solve it with a standard LP-solver, doing so would be slow. As the counting factors and the higher order marginals connecting them form a tree, it is possible to devise a message passing algorithm optimizing (2) exactly. First, this implies that the linear programming relaxation for (2) is exact, as message passing amounts to optimizing the Lagrangian dual of this same relaxation. Second, marginals do not need to be held explicitly, holding messages is enough. The size of all messages equals the size of all counting factors, hence giving again subquadratic space complexity.

Message passing for (2) is detailed in Algorithm 1. It proceeds by first computing *up messages* from adjacent fine subsets to coarser subsets (i.e. going up the tree in Figure 3) and afterwards computing *down messages* from coarse subsets to their equipartition (i.e. going down the tree in Figure 3). Messages reparametrize costs of counting and higher order factors.

Reparametrization Let indices $i < j < l$ be given, where $[i : j]$, $[j+1 : l]$ and $[i : l]$ are subsets generated by the recursive partitioning. Let messages $\phi^{\leftarrow}_{i:j:l}, \phi^{\rightarrow}_{i:j:l}, \phi^{\uparrow}_{i:j:l}$ correspond to constraints (3), (4) and (5) respectively. Messages

ϕ act on (reparametrize) costs θ as

$$\theta^{\phi}_{i:j:l}(x_{i:j}, x_{j+1:l}) \mathrel{+}= \phi^{\leftarrow}_{i:j:l}(x_{i:j}) \qquad \theta^{\phi}_{i:j}(x_{i:j}) \mathrel{-}= \phi^{\leftarrow}_{i:j:l}(x_{i:j}) \tag{6}$$

$$\theta^{\phi}_{i:j:l}(x_{i:j}, x_{j+1:l}) \mathrel{+}= \phi^{\rightarrow}_{i:j:l}(x_{j+1:l}) \qquad \theta^{\phi}_{j+1:l}(x_{j+1:l}) \mathrel{-}= \phi^{\rightarrow}_{i:j:l}(x_{j+1:l}) \tag{7}$$

$$\theta^{\phi}_{i:j:l}(x_{i:j}, x_{j+1:l}) \mathrel{+}= \phi^{\uparrow}_{i:l}(x_{i:l}) \qquad \theta^{\phi}_{i:l}(x_{i:l}) \mathrel{-}= \phi^{\uparrow}_{i:j:l}(x_{i:l}) \tag{8}$$

Algorithm 1: Message passing for one-dimensional discrete tomography

Up messages:
for $[i,j] \cup [j+1,l] = [i,l] \in \Pi$ *in ascending order* **do**
  $\phi^{\leftarrow}_{i:j:l} = \theta_{i:j}$
  $\phi^{\rightarrow}_{i:j:l} = \theta_{j+1:l}$ `// Send messages to higher order counting factor` $\theta^{\phi}_{i:j:l}$
  $\phi^{\uparrow}_{i:j:l}(x_{i:l}) = \min_{x_{i:j}, x_{j+1:l} \sim x_{i:l}} \theta^{\phi}_{i:j:l}(x_{i:j}, x_{j+1:l})$ `// Send message from higher order to counting factor` $\theta^{\phi}_{i:l}$
end
$x^*_{1:n} \in \arg\min_{x_{1:n} \in \mathcal{X}_{1:n}} \theta_{1:n}(x_{1:n})$; `// optimum of top counting factor`
Down messages:
for $[i,j] \cup [j+1,l] = [i,l] \in \Pi$ *in descending order* **do**
  $\phi^{\uparrow}_{i:j:l} = \theta_{i:j:l}$; `// Send message to higher order factor`
  $x^*_{i:j}, x^*_{j+1:l} \in \arg\min_{x_{i:j}, x_{j+1:l} \sim x^*_{i:l}} \theta^{\phi}_{i:j:l}(x_{i:j}, x_{j+1:l})$; `// compute optimal labels`
end

Fast message computation Naively computing one up messages would result in time complexity $O(\ell^5 \cdot n^2)$, which would make the algorithm prohibitively slow. We will describe a fast message computation technique for (3), which uses the structure of the corresponding linear constraints (5) and relies on the latent factorization of $\theta^{\phi}_{i:j:l}$. Specifically, when we fix the endpoints x_{u_i}, x_{u_j} of interval $[i,j]$ and $x_{u_{j+1}}, x_{u_l}$ of $[j+1,l]$, (3) becomes

$$\phi^{\uparrow}_{i:j:l}(x_{i:l}) = \min_{s_{i:j} + s_{j+1:l} = s_{i:l} - x_{u_j} - x_{u_j}} \theta_{u_j, u_{j+1}}(x_{u_j}, x_{u_{j+1}}) + \phi^{\leftarrow}_{i:j}(x_{i:j}) + \phi^{\rightarrow}_{j+1:l}(x_{j+1:l}) \,, \tag{9}$$

Problem (9) is an instance of the min-sum convolution problem: Given $a, b \in \mathbb{R}^n$, compute $c \in \mathbb{R}^{2n-1}$, where $c_i = \min_{j \le i}(a_j + b_{i-j})$. This can be seen by replacing $\phi^{\leftarrow}$ by a, $\phi^{\rightarrow}$ by b and noting that $\theta_{u_j, u_{j+1}}$ is a constant, as x_{u_j} and $x_{u_{j+1}}$ were fixed. For the min-sum convolution problem efficient algorithms [5] were proposed with expected running time $O(n \log(n))$ under the assumption that sorting a and b results in permutations occurring with uniform probability. Problem (9) can be efficiently computed by performing $O(\ell^4)$ min-sum convolutions (one convolution for every choice of endpoints).

Remark 2 (Comparison to [16]). While our approach for solving (2) is inspired by [16], it is notably different: (i) our model includes pairwise potentials forming a chain, while [16] assumes that pairwise potentials do not occur between neighboring subsets. This necessitates to store left and right endpoints in counting factors. (ii) [16] optimizes a different objective: they solve the sum-product version of (2) (i.e. they exchange min by + and + by $\cdot$ in (2)). This allows [16] to use fast Fourier transforms for message computations, instead of the harder min-sum convolution problems.

4 Discrete Tomography Graphical Model

The discrete tomography problem (1) consists of m = #rows(A) distinct one-dimensional subproblems (2). We connect all subproblems (2) via Lagrangian variables into one large problem. This procedure is called dual decomposition, see [15] for an introduction. Specifically, in our discrete tomograpy problems subproblems only share variables $v \in \mathsf{V}$, but not edges $e \in \mathsf{E}$ (shared edges can be handled analoguously). Then for each node $u \in \mathsf{V}$ which participates in the i-th subproblem, we introduce the Lagrangian variable $\lambda_{i,u} \in \mathbb{R}^{|\mathcal{X}_u|}$. The i-th subproblem then consists of solving (2) with the subset of variables U_i, where the unary potentials are the Lagrangian variables $\theta_u = \lambda_{i,u}$. We denote its energy by $E_i(\cdot|\lambda_i)$. The overall problem is

$$\max_{\lambda_1,\dots,\lambda_m} \sum_{i=1}^{m} \min_{x \in \mathcal{X}_{U_i}} E_i(x|\lambda_i) \quad \text{s.t.} \quad \sum_i \lambda_{i,u} \equiv 0 \quad \forall u \in \mathsf{V} \,. \tag{CTG}$$

An exemplary 4×4 model with eight subproblems coming from two projection directions can be seen in Figure 4.

Optimization of relaxation (CTG). To maximize (CTG) we use the bundle solver ConicBundle[2] to find optimal Lagrangian variables λ and Algorithm 1 to find solutions to the one-dimensional subproblems. The bundle method will only give us a dual lower bound to the value of the optimal reconstruction.

Primal solution. To obtain a feasible reconstruction, we solve a reduced problem by excluding labels with high cost: Given dual variables λ_i, let x^* be the optimal solution to the i-th subproblem on variables $U_i \subseteq \mathsf{V}$. For each label $x_u \in \mathcal{X}_u, u \in U_i$, $x_u \neq x^*_u$ we compute the energy $x'^* \in \arg\min_{\{x' \in \mathcal{X}_U : x_u = x'_u\}} E_i(x'|\lambda_i)$ of the minimal reconstruction for subproblem i when the label at u is fixed to x_u (this value can be read off from the reparametrization output by Algorithm 1). Only if the gap $E_i(x'^*|\lambda_i) - E_i(x^*|\lambda_i)$ is smaller than some given threshold, we consider the label x_u. We construct the discrete tomograpy problem on this reduced set of possible labelings and solve the problem with CPLEX [1].

[2] `https://www-user.tu-chemnitz.de/~helmberg/ConicBundle/`

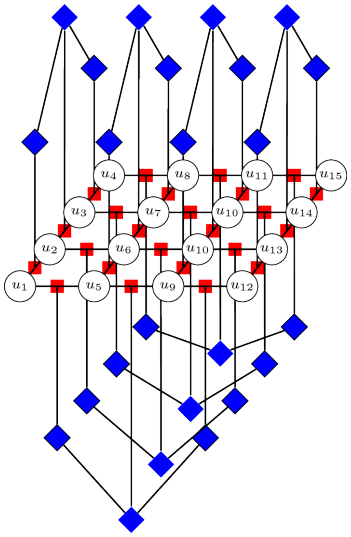

Fig. 4. Illustration of a complete graphical model for the discrete tomography problem with projections along rows and columns of the underlying 4×4 grid. Graphical model consisting of submodel for energy E in (1), indicated by unary (circles) and pairwise (red rectangles) potentials on the 4×4 grid, and counting factors (blue diamond) for the projection constraints $Ax = b$ using the tree construction above the grid (horizontal projection) and below (vertical projection). The higher order potentials are left out for the sake of clarity of presentation.

Comparison to previously used relaxation. It can be shown that the algorithms in [9,19,17,3] use the following relaxation.

$$\begin{aligned}
\min_{\mu \geq 0} \; & \textstyle\sum_{u \in \mathsf{V}} \langle \theta_u, \mu_u \rangle + \sum_{uv \in \mathsf{E}} \langle \theta_{uv}, \mu_{uv} \rangle \\
\text{s.t.} \; & \textstyle\sum_{x_u \in \mathcal{X}_u} \mu_u(x_u) = 1 \quad \forall u \in \mathsf{V} \\
& \textstyle\sum_{x_u \in \mathcal{X}_u} \mu_{uv}(x_u, x_v) = \mu_v(x_v) \quad \forall x_v \in \mathcal{X}_v \qquad \forall uv \in \mathsf{E} \\
& \textstyle\sum_{x_v \in \mathcal{X}_v} \mu_{uv}(x_u, x_v) = \mu_u(x_u) \quad \forall x_u \in \mathcal{X}_u \\
& \textstyle\sum_{u \in \mathsf{V}} A_{iu} \cdot \left(\sum_{x_u \in \mathcal{X}_u} x_u \cdot \mu_u(x_u) \right) = b_i \quad i = 1, \dots, m \,.
\end{aligned} \tag{STD}$$

This relaxation (STD) is the straigtforward generalization of the local polytope relaxation [18] to the discrete tomography problem. The only difference are the linear constraints in the last line of (STD). When specialized to the one-dimensional discrete tomography problem (2), the difference between (STD) and our approach is: for (STD) the tomographic projections are directly enforced through the unary marginals $\mu_u, u \in \mathsf{V}$ instead of enforcing them through the counting factors and higher order ones as we did in Section 3. This more simplistic relaxation (STD) is however less tight.

Proposition 1. *Relaxation* (STD) *is less tight than* (CTG).

Proof. Relaxation (STD) is equivalent to applying it to each tomographic projection separately and then joining every subproblem by Lagrangian variables as we did with our approach above (CTG), see [15, Section 1.6]. Hence, it is enough to show that (STD) is not tight in the one-dimensional case (2). We give a counter-example. Assume $\mathcal{X}_u = \{0, 1\}$ $\forall u \in U$ and we are given Potts pairwise potentials $\theta_{uv}(x_u, x_v) = \begin{cases} 0, & x_u = x_v \\ 1, & x_u \neq x_v \end{cases}$ and zero unary potentials $\theta_u \equiv 0$. Set unary marginals $\mu_u(1) = \frac{b}{|U|}$ and $\mu_u(0) = 1 - \mu_u(1)$ $\forall u \in U$ and pairwise

marginals as $\mu_{uv}(x_u, x_v) = \begin{cases} \mu_u(x_u), & x_u = x_v \\ 0, & x_u \neq x_v \end{cases}$. Such marginals are feasible to (STD), yet give cost 0. On the other hand for e.g. $b = 1$ and $|U| > 1$ there must be at least one label transition, which the Potts potential penalizes with cost 1. □

5 Experiments

Test images. We used 200 randomly generated 32×32 images with three distinct intensity values $\{0, 1, 2\}$, examples of which can be seen in Figure 5. Matrices A for the tomographic projections were constructed as in [13]. For each test image we consider two tomographic problems: (i) measuring along horizontal and vertical directions or (ii) measuring along horizontal, vertical and two diagonal directions (left upper to right lower and left lower to right upper corner). This gives 400 test problems in total. Potentials for energy E in (1) are: unary potentials are zero, while pairwise ones are $\theta_{uv} = |x_u - x_v|$ (that corresponds to TV). Due to integrality of all costs, optimality is ascertained through a duality gap < 1.

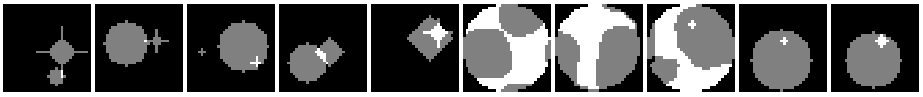

Fig. 5. These images are examples from our testset of 200 images with size 32×32. Black represents value 0, gray 1 and white 2.

Algorithms. We identify our solvers by a prefix **{CTG|STD}** depending on whether (CTG) or (STD) is solved and by a suffix **{CB|relax|BB}** depending on whether ConicBundle, CPLEX [1] or CPLEX with branch and bound enabled was utilized. This gives in total 5 solvers: **CTG_CB**, **CTG_relax**, **CTG_BB**, **STD_relax** and **STD_BB**. We set a timelimit of 1 hour for all algorithms.

Unfortunately, CPLEX cannot solve problems larger than 32×32. When solving the relaxation (CTG), it already consumes multiple GB of memory for 32×32 images. Solving (STD) on the other hand leads to low memory consumption, but CPLEX takes too much time for larger problems (> 1 hour). Hence, to have a baseline, we stick to 32×32 images.

Results. We have proved in Proposition 1 that relaxation (STD) is less tight than our relaxation (CTG). In fact, the first line in Table 2 shows that this occurs 350 times. Furthermore, our tighter relaxation also actually helps in giving optimality certificates. In Table 1 we confirm this numerically: **STD_relax** can provide optimality certificates 53 times, while **CBC_CB** and **CTG_relax** can do so in total 205 times. Interestingly, when using the branch and bound capabilities

of CPLEX, the picture changes and **STD_BB** outperforms **CTG_BB**. This is probably due to the fact that CPLEX can solve the underlying relaxation (STD) much faster than (CTG). We conjecture that the picture will change if the more efficient implementation **CBC_CB** is used as a bounds provider inside a branch and bound solver. This is however outside the scope of our work.

In Figure 6 we give a detailed plot on how much our relaxation (CTG) improved upon (STD).

Also, our relaxation helps in reconstructing the signal. Out of 238 instances, where our heuristic could find an optimal integral solution (third line in Table 2) there were 12 cases, where only our heuristic could do so (second line in Table 2).

	STD relax	STD BB	CTG CB	CTG relax	CTG BB
duality gap "< 1"	53	243	178	154	182
			205		

Table 1. Number of instances where duality gap < 1 (optimality).

	#Instances
(CTG) $>$ (STD) (our relaxation yields **strictly better** lower bound)	350
our heuristic **(only)** found **optimal** integral solution	12
our heuristic found **optimal** integral solution	238

Table 2. Comparison of bounds and primal solutions obtained by (STD) or (CTG).

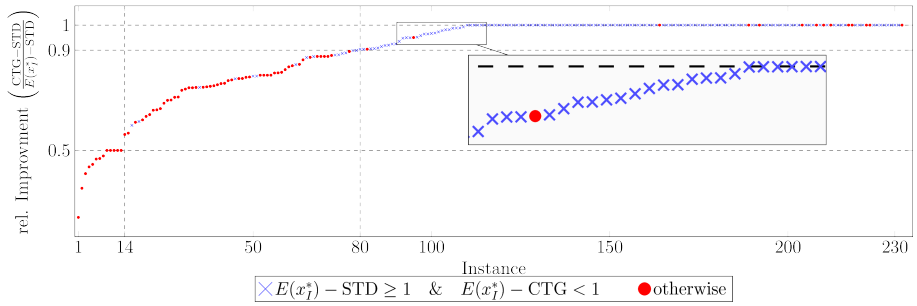

Fig. 6. Comparison of lower bounds (STD) and (CTG). We show relative improvement $^{(\text{CTG})-(\text{STD})}/_{E(x^*)-(\text{STD})}$ on all problems, where we knew by either method the true optimal solution $E(x^*)$, but where (STD) is not tight ((STD) $< E(x^*)$). Lower bound (STD) was computed by **STD_relax**, while (CTG) was computed by **CTG_CB** and **CTG_relax**. A marker close to zero means no improvement and close to one means our relaxation solved the instance exactly. We marked points with a blue cross if only (CTG) but not (STD) achieved a duality gap < 1, i.e. optimality, and used a red circle otherwise. For almost all instance we have an improvement of 0.5 and for more than half of the instances an improvement of 0.9.

6 Conclusion

We have proposed a novel convex relaxation and an accompanying algorithm for the non-binary discrete tomography problem. We have showed theoretically and empirically that our novel relaxation is tighter than the traditionally used relaxation. Solving our new relaxation helps in decoding tomographic reconstructions.

References

1. IBM ILOG CPLEX Optimizer. `http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/`.
2. K. J. Batenburg. An evolutionary algorithm for discrete tomography. *Discr. Appl. Math.*, 151(1):36–54, 2005.
3. K. J. Batenburg. A network flow algorithm for reconstructing binary images from continuous x-rays. *JMIV*, 30(3):231–248, 2008.
4. K. J. Batenburg and J. Sijbers. DART: A practical reconstruction algorithm for discrete tomography. *IEEE TIP*, 20(9):2542–2553, 2011.
5. M. Bussieck, H. Hassler, G. J. Woeginger, and U. T. Zimmermann. Fast algorithms for the maximum convolution problem. *Oper. Res. Let.*, 15:1–5, 1994.
6. B. M. Carvalho, G. T. Herman, S. Matej, C. Salzberg, and E. Vardi. Binary tomography for triplane cardiography. IPMI, 1999.
7. S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhäuser Basel, 2013.
8. E. Gouillart, F. Krzakala, M. Mzard, and L. Zdeborov. Belief-propagation reconstruction for discrete tomography. *Inverse Problems*, 29(3):035003, 2013.
9. J. H. Kappes, S. Petra, C. Schnörr, and M. Zisler. TomoGC: Binary tomography by constrained graph cuts. In *Proc. GCPR*, 2015.
10. S. Keiper, G. Kutyniok, D. G. Lee, and G. E. Pfander. Compressed sensing for finite-valued signals. *ArXiv e-prints*, Sept. 2016.
11. H. Y. Liao and G. T. Herman. Automated estimation of the parameters of Gibbs priors to be used in binary tomography. *Discr. Appl. Math.*, 139(1-3):149–170, 2004.
12. A. Mohammad-Djafari. Gauss-Markov-Potts priors for images in computer tomography resulting to joint optimal reconstruction and segmentation. *International Journal of Tomography & Statistics*, 11(W09):76–92, 2008.
13. S. Roux, H. Leclerc, and F. Hild. Efficient binary tomographic reconstruction. *JMIV*, 49(2):335–351, 2014.
14. T. Schüle, C. Schnörr, S. Weber, and J. Hornegger. Discrete tomography by convex-concave regularization and d.c. programming. *Discr. Appl. Math.*, 151:229–243, 2005.
15. D. Sontag, A. Globerson, and T. Jaakkola. Introduction to dual decomposition for inference. In *Optimization for Machine Learning*. MIT Press, 2011.
16. D. Tarlow, K. Swersky, R. S. Zemel, R. P. Adams, and B. J. Frey. Fast exact inference for recursive cardinality models. In *UAI*, 2012.
17. S. Weber, C. Schnörr, and J. Hornegger. A linear programming relaxation for binary tomography with smoothness priors. In *IWCIA*, 2003.
18. T. Werner. A linear programming approach to max-sum problem: A review. *IEEE TPAMI*, 29(7):1165–1179, 2007.
19. M. Zisler, S. Petra, C. Schnörr, and C. Schnörr. Discrete tomography by continuous multilabeling subject to projection constraints. In *GCPR*, 2016.